Mixed-Initiative Control for Remote Characterization of Hazardous Environments

David J. Bruemmer, Julie L. Marble, Donald D. Dudenhoeffer, Matthew O. Anderson, Mark D. McKay

The Human-System Simulation Laboratory
Idaho National Engineering and Environmental Laboratory
P.O. Box 1625
Idaho Falls, ID 83415-3779, U.S.A.
{bruedj, marbjl, dudedd, matthew, mark}@inel.gov

Abstract

Remote characterization of high radiation environments is a pressing application area where robots can provide benefits in terms of time, cost, safety and quality of data. However, the DOE roadmap for Robotics and Intelligent Machines states that 'usability' may well prove to be the most challenging and yet crucial component of robotic systems for remote characterization and handling of radioactive and hazardous materials. In 2001, the INEEL successfully deployed a teleoperated robotic system coupled with a Gamma Locating and Isotopic Identification Device (RGL&IID) to characterize an area that had been closed to human entry for many years. This paper examines the human-robot dynamic of this teleoperated task and the limitations inherent to the master-slave strategy employed. Next, the paper outlines an innovative, mixed-initiative command and control architecture developed to address these limitations. The resulting, mixed-initiative control architecture retains the human in the loop, but interleaves multiple levels of human intervention into the functioning of a robotic system that can, in turn, scale its own level of initiative to meet whatever level of input is handed down.

1. Introduction

The United States Department of Energy (DOE) continually seeks safer and more cost-effective technologies for use in decontaminating and decommissioning nuclear facilities. To this end, the Deactivation and Decommissioning Focus Area of DOE's Office of Science and Technology sponsors Large-Scale Demonstration and Deployment Projects (LSDDP) to test new technologies. As part of these projects, developers and vendors showcase new products designed to decrease health and safety risks to personnel and the environment, increase productivity, and lower costs.

As part of the FY 2000 and 2001 LSDDP, the Idaho National Engineering and Environmental Laboratory (INEEL) collaborated with the Russian Research and Development Institute of Construction Technology (NIKIMT). This collaboration resulted in the development of the Robotic Gamma Locating and Isotopic Identification Device (RGL&IID) which integrates DOE Robotics Crosscutting (Rbx) technology with NIKIMT Russian gamma locating and isotopic identification technology [1].

While the new robotic solution offered significant improvements in terms of time, cost, worker exposure and the quality of data acquired, the remote nature of this new technology presented new human-robot interaction challenges. Humans were required to enter the building to instrument the environment with cameras and to assist the robot during the execution of the task. Moreover, during the actual deployment, the robot was only allowed to move at very slow speeds due to the limitations of visual feedback to the operator. In answer to these challenges, the INEEL has developed a dynamic autonomy architecture for the same system used in the RGL&IID deployment.

The new approach permits the robot to take initiative to protect itself and the environment. In fact, the human-robot dynamic has changed from a master-slave relationship to that of a mixed team which allows interaction between peers. When compared with the recent RGL&IID technology, the new mixed-initiative system will remove much of the need for prior instrumentation, remove the need for expert operators, reduce the total number of operators, eliminate the need for human exposure and greatly reduce the time needed for preparation and execution of the task.

2. Teleoperated Radiation Survey

Historically at the INEEL, a radiation control technician (RCT) and industrial safety personnel first enter a facility in order to establish accurate conditions for planning purposes. When performing an initial radiation survey, the RCT uses a standard Geiger-Mueller pancake probe to gather radiological information. Once this initial entry has been completed, a video technician may also be required to enter and collect video coverage. Finally a team of sampling technicians is sent into the facility to collect samples used to determine contamination levels and identify which isotopes are present. Typically, this data is then used to aide decontamination and decommissioning (D&D) planning activities (see Figure 1).

Figure 1. Baseline Sample Collection for Laboratory Analysis

2.1. RGL&IID Deployment

To prove that remote, robotic systems could improve on this baseline, the RGL&IID was deployed in July, 2001 at Test Area North (TAN) 616 (see Figure 2). TAN is located at the north end of the INEEL, about 27 miles northeast of the Central Facilities Area. TAN was established in the 1950s by the U.S. Air Force and Atomic Energy Commission Aircraft Nuclear Propulsion Program to support nuclear-powered aircraft research. Upon termination of this research, the area's facilities were converted to support a variety of other DOE research projects. TAN 616 was built in 1954 as a liquid waste treatment facility. As a result of treating thousands of gallons of liquid nuclear processing waste, there are various levels of contamination present in the facility.

Three rooms within TAN 616 were surveyed using the RGL&IID: the Operating Pump Room, the Control Room, and the Pump Room. All of these rooms are filled with process piping and equipment at various levels, which make a manual survey very difficult and time consuming to perform. The intent of this demonstration was to gather empirical data to assess the value of using a remote, robotic system. The metrics considered included reduction in cost, accelerated schedule, improvement in safety, and reliability of data.

Figure 2. TAN 616

2.2. Deployment Results

When compared to baseline assessment methods, the most significant benefit of the RGL&IID deployment was the quality of the results relative to the safety of the workers. Although the RGL&IID deployment did not eliminate the need for workers to enter the contaminated area, it did reduce the need for human exposure. The RGL&IID was compared to the following baseline activities: the initial RCT entry, an entry to collect video, and a final entry to collect sample information. The RGL&IID was able to collect dose information, video coverage, and isotopes present in a single unmanned entry.

Radiation exposure to workers supporting the RGL&IID deployment was cut by more than a factor of 10 over baseline activities. During baseline characterization, workers received 82mRem of radiation exposure. During the deployment of the RGL&IID, workers received 7mRem of radiation exposure. In addition, the RGL&IID provided radiation survey results instantly and the complete facility survey was accomplished in 3 days. It took workers using baseline characterization methods 3 months to accomplish the same results. The baseline activities began in August of 2000 and were not complete until November of 2000. Some of the results from the laboratory analyses were not available until January 2001. The laboratory radiological analysis confirmed the presence of Cs-137, Co-60 and Am-241. This same data was available within minutes after the RG&IILD performed the scan.

The deployment of the RGL&IID did require more workers than the baseline characterization. However, during the baseline sampling activities, six entries with as many as six individuals per entry were made, totaling 60 work hours spent in the contaminated area. During the RGL&IID demonstration, only two technicians and one RCT were required to enter the contaminated facility for a total of 10 work-hours spent in a contaminated area. All others associated with the project were able to complete the objectives from outside the contaminated areas. As a result of workers spending less time in the radiation areas, individuals involved in the RGL&IID deployment received 10 times less radiation dose than workers involved in baseline activities.

In addition, the two technicians and one RCT who did enter the facility during the demonstration did so only to assist the movement of the RGL&IID up and down a flight of stairs and to check air quality prior to entering the facility. These individuals maintained as much distance between themselves and the highest contaminated areas as possible. In contrast, the baseline samplers were required to come in direct contact with the contaminated material in order to collect representative samples.

The financial cost of collecting the radiation measurements using the RGL&IID was about half the cost of the baseline technology. In addition to the benefit of significant cost reductions, this technology also generates significantly more data. For example, whereas the baseline survey included 10 point samples, the RGL&IID collected about 20 scans. Each scan covers as little as one square foot or as much as several square feet and may have as many as 64 point measurements. Altogether the RGL&IID deployment resulted in over 200 point measurements that covered over 100 square feet of wall and floor area. The RGL&IID has the capability of providing 100% coverage if needed.

2.3. Limitations to the Teleoperation Approach

Although the 2001 robotic deployment offered a means to reduce human exposure, it did not fully remove the human from the hazardous environment or make it possible for a single human to control the robot. In fact, the baseline survey required only three people, whereas the RGL&IID required six. If robotic systems are to be truly cost-effective and efficient, this ratio of six humans to one robot must be reduced.

Moreover, the data presented above says nothing about the inherent limitations and risks of teleoperation. Teleoperation requires high-fidelity video, reliable, continuous communication, and costly, potentially dangerous efforts to instrument the environment *a priori*. As a mechanical 'subordinate,' the robot was dependent on continuous, low-level input from a human and was poorly equipped to cope with communication failures or changes in operator workload. In fact, while training within a mock-up facility, operators lost control of the vehicle due to a communication failure. Since the last command received by the robot before communications were lost had been a forward velocity command, the robot continued to navigate across the room and actually ran right through the walls of an adjacent test bed environment. As a result, the robot's control system was immediately changed to have a "watchdog" system that halted the robot once it recognized that communications had failed.

Even so, communication proved to be the limiting factor governing human-robot interaction during the teleoperated deployment. Thick concrete shielding, typical to radiological controls, made it extremely difficult for high-bandwidth communication to support the strictly teleoperated system. As a result, it was necessary for a human to physically place a large antenna directly into the opening of the TAN 616 building. As the robot traveled further from this antenna, the possibility of communication dropouts increased. In fact, operators completely lost contact with the robot at one point during the deployment when the robot traveled out of range. The robot stopped after several seconds once it recognized that communication had been lost. Since the robot was merely a passive tool, it was unable to reorient itself or attempt to reestablish communication. If humans had been unable to enter the environment, the robot would have been lost and unable to complete its task. Fortunately, a human was able to move the antenna slightly further into the doorway of the building and communication was reestablished.

The 2001 RGL&IID deployment required weeks of preparation including training operators in mock-up environments. Early on, these training exercises indicated that cameras positioned on the robot would not be sufficient to support teleoperation. The camera could not see the immediate obstacles surrounding the wheels – the very obstacles that posed the greatest threat. As a result, it was necessary to instrument the environment *a priori* with elevated cameras. These cameras were tethered to allow sufficient bandwidth for high resolution video and were set up in the environment by humans. Human placement of tethered cameras is a common practice in nuclear remote inspections throughout the DOE complex. This drawback to teleoperated approaches is further pronounced by the fact that these cameras must be bagged, resulting in additional contaminated waste once the operation is complete.

Although the cameras were deemed sufficient for the task, operators explained that such a strategy is inherently limiting. The first limitation is that adequate lighting is required to support vision-based teleoperation. Secondly, such cameras are usually unable to provide complete visual coverage. In fact, operators reported blind spots when using the same robotic system and cameras within a different, larger building at the site. In one instance, as the robot rounded a corner and left the visual field of one camera, the last thing the operators saw was the robot begin to tip over. Fortunately, the robot did not tip over and the task was completed successfully. Nonetheless, the incident emphasizes the need for the robot to provide better feedback and, ideally, to be able to take initiative to protect itself in critical situations.

3. Mixed-Initiative Control

Throughout the DOE complex, teleoperated systems have often failed to address the limitations of telepresence inherent to current communication technologies. On the other hand, attempts to build and use autonomous systems have failed to acknowledge the inevitable boundaries to what the robot can perceive, understand, and decide apart from human input. Both approaches have failed to build upon the strengths of the robot and the human working as a cohesive unit. In response to limitations of both approaches, research efforts at the INEEL have developed a novel robotic system that can leverage its own, intrinsic intelligence to support a spectrum of control levels. We submit that rather than conceive of machines as mere tools or, on the other hand, as totally autonomous entities that act without human intervention, it is more effective to consider the machine as part of a dynamic human-machine team. Within this team, each member is invested with agency – the ability to actively and authoritatively take initiative to accomplish task objectives. Within this schema, each member has equal responsibility for performance of the task, but responsibility and authority for particular task elements shifts to the most appropriate member, be it human or machine. For instance, in a remote situation, the robot may be in a much better position than the human to react to the local environment, and consequently, the robot may take the leadership role regarding navigation. As leader, the robot can then “veto” dangerous human commands to avoid running into obstacles or tipping itself over.

The resulting robotics system, pictured in Fig. 3., including hardware, software, and interface components, can slide between roles of ‘subordinate,’ ‘equal’ and ‘leader.’ The ability of the robot to change its level of autonomy on the fly supports changing communication, cognitive, perceptual and action capabilities of the user and robot. With the new system, communications dropouts no longer result in the robot stopping dead in its tracks or, worse, continuing rampant until it has recognized that communications have failed. Instead, the robot may shift into a fully autonomous mode.

For this system to meet its goals, we must provide robust mechanisms which allow the robot to protect itself and the environment. To do so we fuse a variety of range sensor information including inertial sensors, compass, wheel encoders, laser range finders, computer vision, thermal camera, infrared break beams, tilt sensors, bump sensors, sonar, and others. The robot does not assume that these sensors are working correctly, but rather continuously evaluates its own perceptual capabilities and behavior. Novel sensor-suites and fusion algorithms enhance capabilities for sensing, interpreting, and "understanding" environmental features. Also, a great deal of work has focused on providing situation awareness to the user that can appropriately support the current level of interaction. With the new system we are not limited to visual feedback. Instead, the robot is able to abstract information about the environment at many levels including terse textual descriptions of the robot’s local surroundings.

Given the desire to employ robots in hazardous, critical environments, the ability to shift a robot in and out of the leadership role presents a conundrum. The user comes to rely on the self-protective capabilities of the robot and yet, at times, must override them to accomplish a critical mission. For instance, when faced with an unknown obstacle that fully obstructs the path, the user may shift the robot out of the leadership responsibility for navigation, but grant the robot the “right” to refuse human commands when the physical resistance to motion is beyond a certain threshold. This allows the human to attempt to push the obstacle out of the way without exerting dangerously high force on the robot. For other tasks, the user may need to drive the robot to where it is touching an obstacle in order to take a sample. The user can curtail the robot’s collision avoidance initiative and yet customize a “last resort” channel of initiative based on bump sensors and short-range infrared break beams.

Ideally, we need control systems that allow the user to configure the autonomy of the robot on the fly, activating “channels of initiative” that crosscut broad categories. The roles of each team member are bounded by a complex and changing web of capabilities and limitations to which each member must adapt and respond. The ability of the human to develop accurate understanding of robot behavior is essential if this adaptive role switching is to work effectively. One of the most fascinating areas of future work is the need for the robot to be imbued with an ability to understand and predict human behavior.

3.1. Theory of Robot Behavior

The need for both human and robot to predict and understand one another's actions presents a daunting challenge. For each level of robot initiative, the user must develop a unique set of expectations regarding how the robot behaves, that is, an understanding or theory of the system's behavior, here after referred to as a theory of robot behavior (TORB). By TORB we mean that the human operator is able to quickly and accurately predict:

1. Actions the robot will take in response to stimuli from the environment and other team members;
2. The outcome of the cumulative set of actions.

In our research we are not concerned with developing a formal model of robot cognition, but rather require that the human understand and predict the emergent actions of the robot, with or without an accurate notion of how intelligent processing gives rise to the resulting behavior. When a human team member is faced with a robot that can orchestrate task elements, the critical issue will not be how the robot or machine "reasons," but rather whether the human team members can accurately predict robotic responses and understand how cumulative actions and responses converge to fulfill task objectives.

Many applications require the human to quickly develop an adequate TORB. One way to make this possible is to leverage the knowledge humans already possess about human behavior and other animate objects, such as pets or even video games, within our daily sphere of influence. For example, projects with humanoids and robot dogs have explored the ways in which modeling emotion in various ways can help (or hinder) the ability of a human to effectively formulate a TORB [2],[3].

Regardless of how it is formed, an effective TORB allows humans to recognize and complement the initiative taken by robots as they operate under different levels of autonomy. The ability to predict and exploit the robot's initiative will build operator proficiency and trust. The development of a theory of robot behavior will also allow the user to switch between and configure the robot's levels of initiative to suit the needs and components of the task at hand.

3.2. Theory of Human Behavior

Just as the human develops a theory of the robot's behavior, the robot must be able to understand and predict the human members of the team in order to adapt to their needs. This is not to say that machines must possess complex mental models or be able to discern our intentions. Rather, it is necessary to raise the level of interaction between the human and robot based upon readily available, non-intrusive workload cues emanating from the operator. The robot's theory of human behavior may be a rule set at a very simple level, or it may be a learned expectation developed through practiced evolutions with its human counterpart. The robot must possess some means to infer the need for intervention. Currently, accurate and non-intrusive collection of these cues is difficult at best, and those measures that have been used are unreliable at worst [4].

The answer to this dilemma is to reduce the human signals down to a prescribed set of channels, which are available as an integral part of the interaction of the human with the machine, and which the machine can use to configure its behavior and level of initiative. Interaction between the robot and human may be through direct communications (verbal, gesture, touch, radio communications link) or indirect observation (physically struggling, erratic behavior, unexpected procedural deviation). Interaction may also be triggered by the observation of environmental factors (rising radiation levels, the approach of additional humans, etc.). The robot's expectations must allow it to recognize human limitations and anticipate human needs without second-guessing the human's every move. When robots do intervene with their human counterparts, the human's TORB must be able to explain why the robot has stepped in and what this shift in control means for the task at hand.

3.3. Dynamic Role Changing

The benefits of allowing the team members to change roles within the team significantly increases team flexibility and reliability in task performance. However, if the interface and human-robot system are not designed in accordance with critical principles of human factors in mind, dynamic role changing may result in mode confusion, loss of operator situation awareness, loss of operator confidence in assuming supervisory control, and degraded and potentially catastrophic performance [5]. Systematic human-centered design is necessary to insure that the robot autonomy conforms to the ways in which humans assign and manage tasks.

Appropriate feedback is required when roles and levels of initiative change. Failure to inform the operator when the robot has overridden commands will lead to distrust of the system, unless the behavior is beneath the level of operator concern. This phenomenon has been studied within the airline industry with pilots and the automatic pilot mode of operation. [6]. Feedback from the robot should not only include the mode change, but also an indication of the reason for the change. For optimal performance of the team, the human must be able to develop expectations regarding when and why the robot will be motivated to initiate a new level of initiative. In order for the human's theory of system behavior to comprehend and exploit robot initiative, the robot's autonomy should be structured hierarchically such that at any given time, the user will know the bounds on what initiative the robot can take. Consequently, the INEEL has developed a control system that supports four clearly distinct levels of human intervention.

4. Levels of Human Intervention

Within the last five years, researchers have begun in earnest to examine the possibility for robots to support multiple levels of user intervention. Much of this work has focused on providing the robot with the ability to accept high level verbal, graphical, and gesture-based commands [7], [8], [9]. Others have implemented robots that understand the limitations of their autonomous capabilities and can query the user for appropriate assistance [10], [11]. Goodrich et al. [12] have performed experiments which involve comparing the performance of human-robot pairs using different modes of human intervention.

However, very little work has emphasized true peer to peer interactions where the robot is actually able to shift modes of autonomy as well as the user. Sholtz [13] discusses the need for this kind of peer-peer interaction, and provides categories of human intervention including

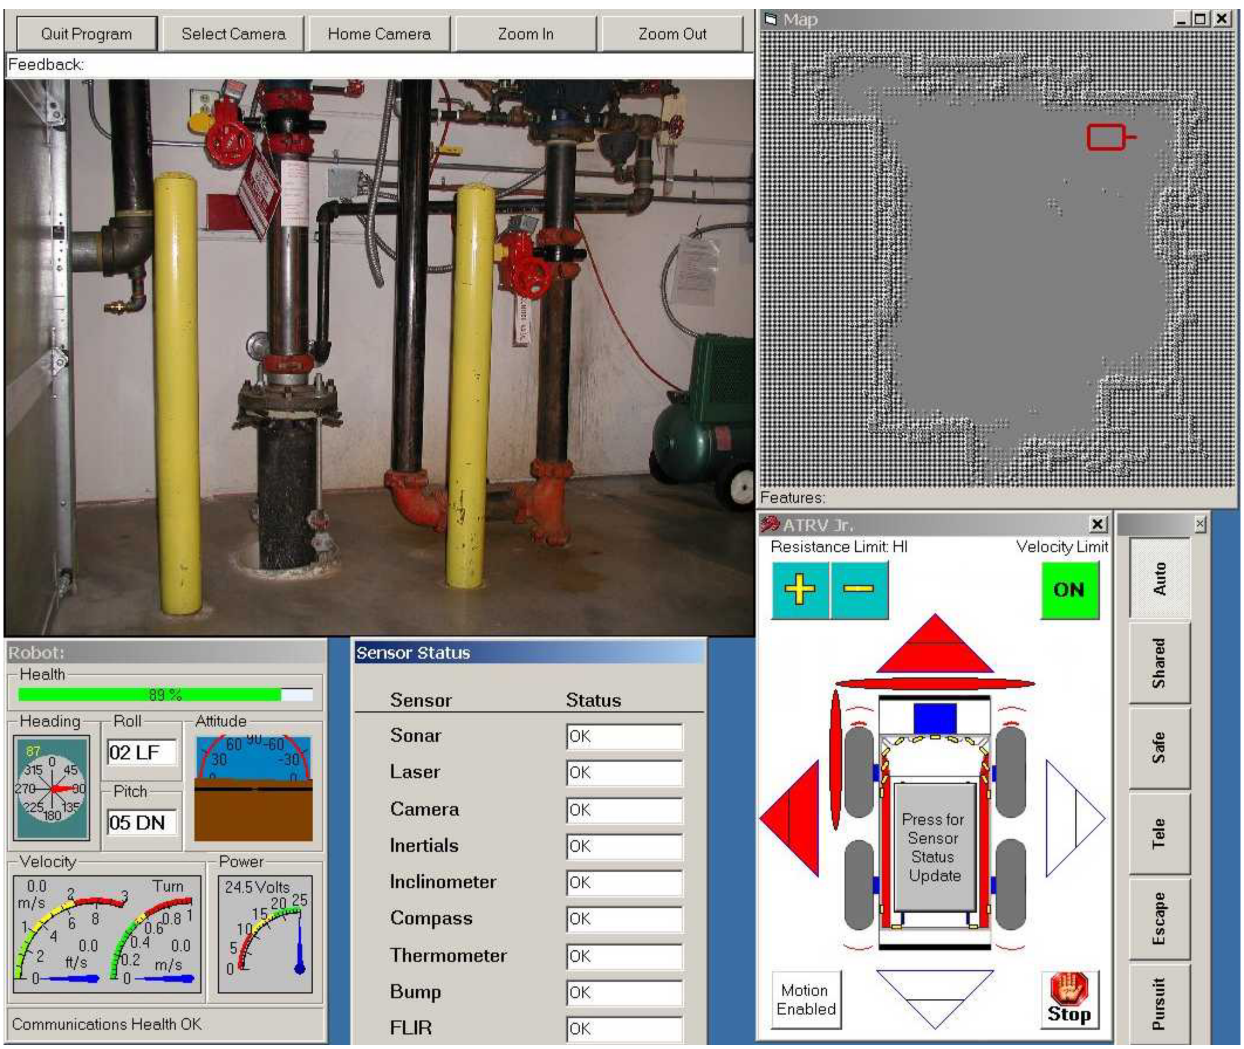

Figure 4: Current interface used for mixed-initiative control of the robot

supervisory, peer to peer and mechanical interaction (e.g. teleoperator). Our research to date has developed a control architecture that spans these categories, supporting the following modes of remote intervention:

1. Teleoperation
2. Safe Mode
3. Shared Control
4. Full Autonomy

For each of these levels of autonomy, perceptual data is fused into a specialized interface (shown in figure 4) that provides the user with abstracted auditory, graphical and textual representations of the environment and task that are appropriate for the current mode. Currently, this interface is used on a touch screen tablet PC made by Fujitsu Corp.. Within this interface, blockages are shown as red ovals and resistance to motion is shown as arcs emanating from the wheels. The robot relays a great deal of synthesized, high-level information (including suggestions and requests for help) to the user in a textual form using the feedback textbox within the image window. Also note that the robot provides textual reports on environmental features at the bottom of the map window and reports on communications status at the bottom of the robot status window. The robot status window provides a variety of information about the status of the robot including pitch and roll, power, heading, speed and a fusion of this information into a single measurement of "health."

The user can move the robot by touching the arrows or may use a joystick or other game controller. It is possible to pan and tilt the camera automatically by touching regions of the visual image. Currently, we are still working to integrate the on-the-fly mapping capabilities with the interface shown in figure 4. As we continue this task, the interface will allow a number of autonomous tasks (e.g. searching a specified region or going to a goal location) to be issued by interacting with the map itself.

4.1. Teleoperation

We have taken the interaction substrate used in previous INEEL teleoperated robotic systems and revamped it through feedback from people who have deployed such systems. Within teleoperation mode, the user has full, continuous control of the robot at a low level. The robot takes no initiative except to stop once it recognizes that communications have failed.

4.2. Safe Mode

Within safe mode, the user directs the movements of the robot, but the robot takes initiative to protect itself. In doing so, this mode allows the user to issue motion commands with impunity, greatly accelerating the speed and confidence with which the user can accomplish remote tasks. The robot assesses its own status and surrounding environment to decide whether commands are safe. For example, the robot has excellent perception of the environment and will stop its moti on just before a collision, placing minimal limits on the user to take the robot's immediate surroundings into account. The robot also continuously assesses the validity of its diverse sensor readings and communication capabilities. The robot will refuse to undertake a task if it does not have the ability (i.e., sufficient power or perceptual resources) to safely accomplish it.

4.3. Shared Control

The robot takes the initiative to choose its own path, responds autonomously to the environment, and works to accomplish local objectives. However, this initiative is primarily reactive rather than deliberative. In terms of navigation, the robot responds only to its local (~ 6-10 meter radius), sensed environment. Although the robot handles the low level navigation and obstacle avoidance, the user supplies intermittent input, often at the robot's request, to guide the robot in general directions. The problem of deciding how and when the robot should ask for help has been a major line of HRI enquiry and will be a major issue in our upcoming human subject experiments.

4.4. Full Autonomy

The robot performs global path planning to select its own routes, requiring no user input except high-level tasking such as "follow that target" or "search this area" specified by drawing a circle around a given area on the map created by the robot. This map is built on the fly and uses frontier-based exploration and localization to perform searches over large areas including multiple rooms and corridors. The user interacts with the map to specify tasks and can guide the robot and infuse knowledge at an abstract level by selecting areas of interest and identifying sensed environmental features, which then become included within the map.

These levels of operator intervention can greatly improve on the opportunities provided to the operators of a strictly teleoperated system such as the one used in the RGL&IID deployment. The human user can switch between these modes to cope with different components of the task. For instance, when a user wishes to move into a new room s/he simply points the robot at a door and then allows the robot to guide itself through the doorway – a task that reportedly took teleoperators many minutes of trial and error.

The latest development, and perhaps the most innovative aspect of our project to date, is that we have imparted a "theory of human behavior" within the robot's intrinsic intelligence, which allows the robot to assess human performance. Before we implemented this theory of human behavior, the robot was already able to use its knowledge of the environment and its own proprioception to take initiative and refuse to accept dangerous commands. However, the level of robot initiative was always controlled by the human. The “theory of human behavior” allows the robot to switch modes when the robot recognizes that the human is performing very poorly. This theory of human behavior is based primarily on the frequency of human input and the number and kind of dangerous commands issued by the user. For instance, if the human has repeatedly placed the robot or the environment in danger, or if the human has been unsuccessful in extricating a robot from a cluttered area, the robot will step in and take over from the operator. Although the human can ultimately override this capability, it provides a means for true peer-peer interaction.

5. Conclusions

The INEEL is currently exploring new ground in the area of human interaction with robots. The motivation for our work is the development of flexible human-robot teams to support the performance of tasks within human-hazardous environments. Most importantly, the prevailing trend that forces humans to adapt to the limits of inflexible technology will give way to interface technologies that adapt to our needs, enabling a new era of human-machine interaction.

For a robotic system to gracefully accept a full spectrum of intervention possibilities, interaction issues cannot be handled merely as augmentations to a control system. Instead, opportunities for operator intervention must be incorporated as an integral part of the robot’s intrinsic intelligence. The robot must be imbued with the ability to accept different levels and frequencies of intervention. Moreover, for autonomous capabilities to evolve, the robot must be able to recognize when help is needed from an operator and/or other robot and learn from these interactions.

Utilizing a robot equipped with robust sensors and intelligence, we have developed a human-robot control system and associated interfaces that promote mutual-initiative between the human operator and the robot. The robot is often able to make better judgments about its environment (i.e., local navigation) than distal human controllers. Consequently, we have created modes of control where the robot monitors human command input and infers the need to supplement or override human action. The robot has the power to refuse to undertake commands from the user that are deemed by the robot to pose a threat to itself or its environment. This engenders a host of new questions, especially in regard to how an autonomous and mobile robot can infer intervention points. Within our implementation, human error loses much of its sting because the robot is able to countermand dangerous commands. At the same time, we have provided means for the human to override robot initiative and to configure the robotic initiative for specific tasks and environments. In this way, the human and robot become true team partners who can support and compensate for one another to adapt to new challenges.

6. References

[1] Anderson, M., Conner, C., Daniel, V., McKay, M., Yancey, N., Demonstration of the Robotic Gamma Locating and Isotopic Identification Device, In Proceedings from the American Nuclear Society Spectrum 2002, Reno, NV, August 2002.

[2] Brooks, R., Brazeal, C., Marjanovic, M., Scassellati, B., & Williamson, M. “The Cog Project: Building a Humanoid Robot.” In *Computation for Metaphors, Analogy, and Agents, Springer Lecture Notes in Computer Science, 1562*, C. Nehaniv, Ed., Springer-Verlag, Berlin, 1998.

[3] Brooks, R. A., *Flesh and Machines*, New York: Pantheon Books, 2002, pp. 51-52.

[4] Wickens, C. D., *Engineering Psychology and Human Performance,* (2nd ed.), NY, NY: Harpers Collins Publishers, Inc. 1992.

[5] M. A. Goodrich and E. R. Boer. Designing Human-Centered Automation: Tradeoffs in Collision Avoidance System Design. *IEEE Transactions on Intelligent Transporation Systems*, v.1,no.1, March 2000.

[6] Abbott, K. A., S. M. Slotte, D. K. Stimson, “Federal Aviation Administration Human Factors Team Report on: The Interfaces Between Flightcrews and Modern Flight Deck Systems,” Federal Aviation Administration, Washington, DC, Tech. Rep. June, 1996. [Online]. Available: http://www.faa.gov/avr/afs/interfac.pdf, pp. D1-D3.

[7] D. Perzanowski, A.C. Schultz, W. Adams, M. Skubic, M. Abramson, M. Bugajska, E. Marsh, J. G. Trafton, and D. Brock, Communicating with Teams of Cooperative Robots, In Proceedings from the *2002 NRL Workshop on Multi-Robot Systems*, Washington, D. C., March, 2002.

[8] Kortenkamp, D., Huber , E., and Bonasso, R. P., Recognizing and interpreting gestures on a mobile robot. In Proc. of AAAI 1996. Portland, OR, August, 1996.

[9] Voyles, R., and Khosla , P., Tactile gestures for human-robot interaction. In Proc. of IEEE/RSJ Conf. On Intelligent Robots and Systems, Volume 3, 7-13. 1995.

[10] Fong, T., Thorpe, C., Baur, C., Robot as Partner: Vehicle Teleoperation with Collaborative Control, In Proceedings from the *2002 NRL Workshop on Multi-Robot Systems*, Washington, D. C., March, 2002.

[11] Murphy, R. R., and Rogers, E. 1996. Cooperative assistance for remote robot supervision. *Presence* 5(2): 224-240.

[12] Goodrich, M. A. and Olsen, D. R., Crandall, J. W. and Palmer, T. J. Experiments in Adjustable Autonomy, In *IJCAI 2001 Workshop on Autonomy Delegation and Control*, Seattle WA, August, 2001.

[13] Scholtz, J., Human-Robot Interactions: Creating Synergistic Cyber Forces, In Proceedings from the *2002 NRL Workshop on Multi-Robot Systems*, Washington, D. C., March, 2002.